Cache Assist in Hard Drives

SNIA Forward Looking Information Disclosure Statement

This SNIA presentation as part of the industry EPA ENERGYSTAR Data Center Storage Stakeholders Meeting November 18 2015 may include timetables, roadmaps, new technologies entering the mainstream, predictions, estimates or other information that might be considered forward-looking. While these forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual timeframes and results to differ materially. Readers are cautioned not to place undue reliance on these forward-looking statements, which reflect our opinions and best effort planning only as of the date of this presentation. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forward-looking statements in light of new information or future events. Throughout the discussion in the delivery of this presentation, we will attempt to present some important factors relating to the topic that may affect our estimates and predictions.

What is Cache Assist

SNIA Green Storage Initiative

- Any form of HDD caching that enhances performance without negatively impacting data integrity.
- Until several years ago we were here on enterprise products ….
  - Multi-segmented DRAM. Varying sizes. 128MiB not unusual
  - Read sophistication was limited to cache and “read look ahead”.
  - Writes can also utilize DRAM. Two options on how the drive informs the host of write status.
    - Write Cache Enabled – Faster Performance – Risky if there is a power hiccup.
    - Write Cache Disabled – Slower Performance - No risk of data loss due to power hiccups
- Implication – since vast majority of enterprise customers run with HDD write caching off, performance at the drive was to a large degree closely coupled to raw performance, i.e. sustained disc data rates, seek performance profiles, and RPM.

Change drivers

What is the typical workload for servers by application?

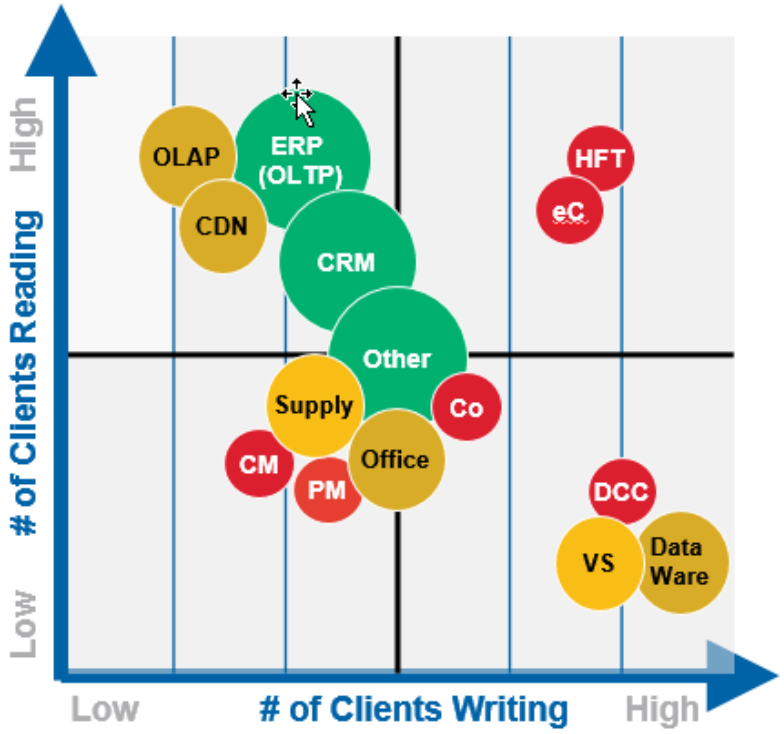

- Small overall IT Spend (<$50B)
- Medium overall IT Spend ($50 to $100B)
- Large overall IT Spend (>$100B)
- Size of bubble indicates amount of IT spend

Takeaways:

- Majority of IT SW spend is on read intensive applications
- Read intensive applications will be challenged as the number of clients reading data increases
  (Integration of ERP, CRM, and OLAP)

CRM: Customer Relationship Management
ERP: Enterprise Resource Management (OLTP)
OLAP: Data Mining & Business Intelligence
CDN: Content Delivery Networks, Broadcast
Office: General Office Suites
Supply: Supply Chain Management
VS: Video Surveillance
eC: eCommerce

HFT: High Frequency Trading
Data Ware: Data Warehousing
CM: Content Management
PM: Project Management
DCC: Digital Content Creation
Co: Web Conferencing, Conferencing, Social
Other: Other

What is Changing

- Starting with the appearance of 4K disc formats, NVC was used for torn write protection and read-modify-write in 512e configurations

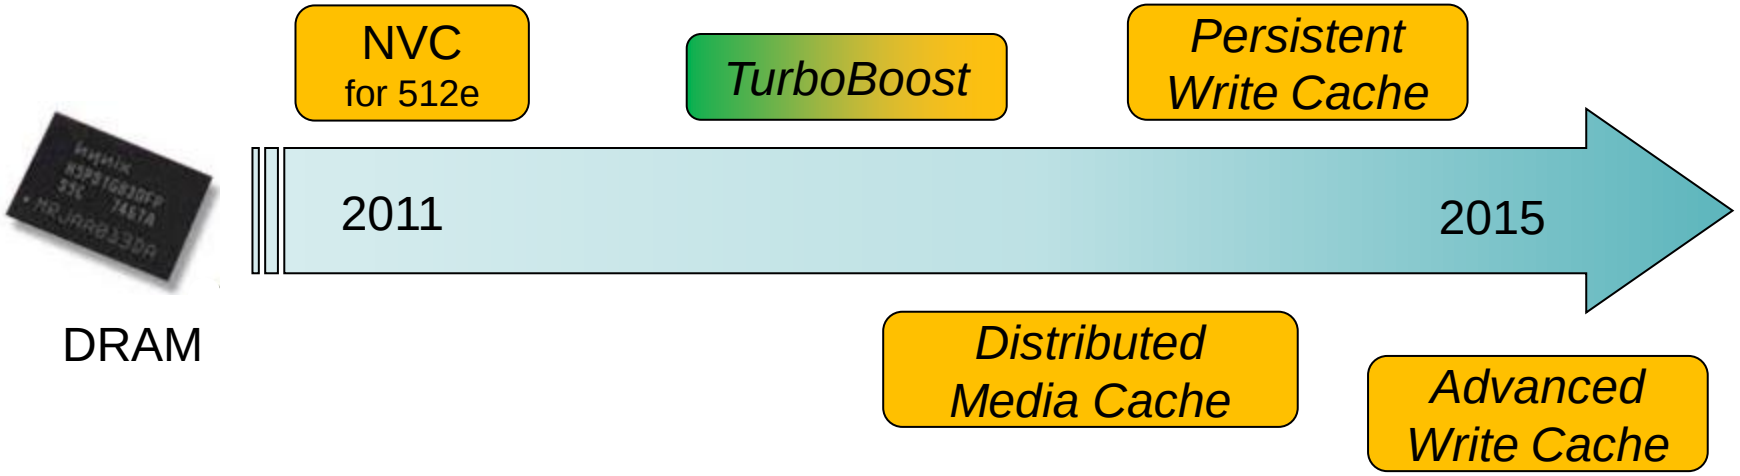

Improved Performance in localized workloads
Write cache enabled performance with the data security of write cache disabled

Various ways to implement

SNIA Green Storage Initiative

Cache Technology Name	Physical	Effect on Reads	Effect on Writes
Seagate TurboBoost	Large NAND for reads. NVC for writes	Large size and sophistication results in greatly improve response times for localized IO	Improved write performance. Visible in classic random writes
HGST Distributed Media Cache	Distributed caching regions on discs	None	Improved write performance. Visible in classic random writes
Toshiba Persistent Write Cache	NOR Flash & perhaps some DMC	None	Improved write performance. Visible in classic random writes
Seagate Advanced Write Cache	?	None	Improved write performance. Visible in classic random writes

What does it all mean ?

SNIA Green Storage Initiative

- Cache assist technologies can result in decoupling of performance from what might be expected given classic specs of SDR, seek time, etc
  - The performance changes in Random Writes and HotBand can be significant.
  - Power changes associated with performance increases might not be as expected in years past.
- Some technologies are separately available features, some are not
  - Currently available read cache implementations i.e. TurboBoost, are distinct SKU's.
  - New write caching technologies are embedded in the base product and very likely cannot be disabled.
    - NOR based NVC sizes are frequently not specified in data sheets.
    - Media caches have no specified size and are not spec'd
- Small amounts of NVC on 512n drives for performance is not out of the question.
- Interaction of drive write caches (that cannot be disabled) with controller caches.

Hot Band and Random Results

	Seagate 1.8TB 10K	Seagate 1.8TB 10K TurboBoost	Seagate 8TB Nearline
SDR	229	232	237
Read iops	186	183	117
Write Iops	342	399	362
Hot Band IOPSpw	19.8	35.3	9.2
Hot Band IOPS	1,800.1	3,227.3	1,052.1
Hot Band Power	90.8	91.4	114.8
Ran Write IOPSpw	9.4	13.0	9.3
Ran Write IOPS	866.0	1,165.7	990.9
Ran Write Power	92.4	89.8	106.5
Ran Read IOPSpw	13.3	13.5	4.7
Ran Read IOPS	1,201.0	1,209.9	519.8
Ran Read Power	90.5	89.7	111.2

Sequential Results

SNIA Green Storage Initiative

	Seagate 1.8TB 10K	Seagate 1.8TB 10K TurboBoost	Seagate 8TB Nearline
SDR	229	232	237
Read iops	186	183	117
Write Iops	342	399	362
Seq Write MiB/sec/w	2.8	3.7	1.6
Seq Write MiB/sec	234.6	304.4	177.0
Seq Write Power	85.2	82.5	109.8
Seq Read MiB/sec/w	14.1	14.2	10.7
Seq Read MiB/sec	1,167.0	1,170.8	1,173.9
Seq Read Power	82.6	82.2	109.47

Configuration for this and prior slide
RAID 10 using 6 drives.
Streaming done with 2 streams